# On the automated verification of web applications with embedded SQL

**Shachar Itzhaky[1], Tomer Kotek[2], Noam Rinetzky[3], Mooly Sagiv[3], Orr Tamir[3], Helmut Veith[2], and Florian Zuleger[2]**

1 **Massachusetts Institute of Technology**
2 **Vienna University of Technology**
3 **Tel Aviv University**

arXiv:1610.02101v1 [cs.LO] 6 Oct 2016

**Abstract**

A large number of web applications is based on a relational database together with a program, typically a script, that enables the user to interact with the database through embedded SQL queries and commands. In this paper, we introduce a method for formal automated verification of such systems which connects database theory to mainstream program analysis. We identify a fragment of SQL which captures the behavior of the queries in our case studies, is algorithmically decidable, and facilitates the construction of weakest preconditions. Thus, we can integrate the analysis of SQL queries into a program analysis tool chain. To this end, we implement a new decision procedure for the SQL fragment that we introduce. We demonstrate practical applicability of our results with three case studies, a web administrator, a simple firewall, and a conference management system.

**1998 ACM Subject Classification** D.3.2 Database Management Languages, F.3.1 Specifying and Verifying and Reasoning about Programs

**Keywords and phrases** SQL; Scripting language; Web services; Program verification; Two-variable fragment of First Order logic; Decidability; Reasoning

**Digital Object Identifier** 10.4230/LIPIcs.xxx.yyy.p

## 1 Introduction

Web applications are often written in a scripting language such as PHP and store their data in a relational database which they access using SQL queries and data-manipulating commands [36]. This combination facilitates fast development of web applications, which exploit the reliability and efficiency of the underlying database engine and use the flexibility of the script language to interact with the user. While the database engine is typically a mature software product with few if any severe errors, the script with the embedded SQL statements does not meet the same standards of quality.

With a few exceptions (such as [15, 19]) the systematic analysis of programs with embedded-SQL statements has been a blind spot in both the database and the computer-aided verification community. The verification community has mostly studied the analysis of programs which fall into two classes: programs with (i) numeric variables and complex control structure, (ii) complex pointer structures and objects; however, the modeling of data and their relationships has not received the same attention. Research in the database community on the other hand has traditionally focused on correct design of databases rather than correct use of databases.

This paper lays the ground for an interdisciplinary methodology which extends the realm of program analysis to programs with embedded SQL. Since the seminal papers of Hoare, the

© Shachar Itzhaky, Tomer Kotek, Noam Rinetzky, Mooly Sagiv, Orr Tamir, Helmut Veith, and Florian
licensed under Creative Commons License CC-BY
Conference title on which this volume is based on.
Editors: Billy Editor and Bill Editors; pp. 1–25
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

first step for developing program analysis techniques is a precise mathematical framework for defining programming semantics and correctness. In this paper we develop a Hoare logic for a practically useful but simple fragment of SQL, called SmpSQL, and a simple scripting language, called SmpSL, which has access to SmpSQL statements. Specifically, we describe a decidable logic for formulating specifications and develop a weakest precondition calculus for SmpSL programs; thus our Hoare logic allows to automatically discharge verification conditions. When analyzing SmpSL programs, we treat SQL as a black box library whose semantics is given by database theory. Thus we achieve verification results relative to the correctness of the underlying database engine.

We recall from Codd's theorem [13] that the core of SQL is equivalent in expressive power to first-order logic FO. Thus, it follows from Trakhtenbrot's theorem [34] that it is undecidable whether an SQL query guarantees a given post condition. We have therefore chosen our SQL fragment SmpSQL such that it captures an interesting class of SQL commands, but corresponds to a decidable fragment of first-order logic, namely $\mathrm{FO}^2_{\mathrm{BD}}$, the restriction of first-order logic in which all variables aside from two range over fixed finite domains called *bounded domains*. The decidability of the finite satisfiability problem of $\mathrm{FO}^2_{\mathrm{BD}}$ follows from that of $\mathrm{FO}^2$, the fragment of first-order logic which uses only two variables. Although the decidability of $\mathrm{FO}^2$ was shown by Mortimer [29] and a complexity-wise tight decision procedure was later described by Grädel, Kolaitis and Vardi [21], we provide the first efficient implementation of finite satisfiability of $\mathrm{FO}^2$.

We illustrate our methodology on the example of a simple web administration tool based on [22]. The PANDA web administrator is a simple public domain web administration tool written in PHP. We describe in Section 2 how the core mailing-list administration functionality falls into the scope of SmpSL. We formulate a specification consisting of a database invariant and pre- and postconditions. Our framework allows us to automatically check the correctness of such specifications using our own $\mathrm{FO}^2_{\mathrm{BD}}$ reasoning tool.

### Main contributions

1. We define SmpSQL, an SQL fragment which is contained in $\mathrm{FO}^2_{\mathrm{BD}}$.
2. We define a a simple imperative script language SmpSL with embedded SmpSQL statements.
3. We give a construction for weakest preconditions in $\mathrm{FO}^2_{\mathrm{BD}}$ for SmpSL.
4. We implemented the weakest precondition computation for SmpSL.
5. We implemented a decision procedure for $\mathrm{FO}^2_{\mathrm{BD}}$. The procedure is based on the decidability and NEXPTIME completeness result for $\mathrm{FO}^2$ by [21], but we use a more involved algorithm which reduces the problem to a SAT solver and is optimized for performance.

We evaluate our methodology on three applications: a web administrator, a simple firewall, and a conference management system. We compared our tool with Z3 [14], currently the most advanced general-purpose SMT solver with (limited) support for quantifiers. In general, our tool performs better than Z3 in several examples for checking the validity of verification conditions of SmpSL programs. However, our tool and Z3 have complementary advantages: Z3 does well for unsatisfiable instances while our tool performs better on satisfiable instances. We performed large experiments with custom-made blown up versions of the web administrator and the firewall examples, which suggest that our tool scales well. Moreover, we tested the scalability of our approach by comparing of our underlying $\mathrm{FO}^2$ solver with three solvers on a set of benchmarks we assembled inspired by combinatorial problems. The solvers we tested against are Z3, the SMT solver CVC4 [3], and the model checker Nitpick [7]. Our solver outperformed each of these solvers on some of the benchmarks.

## 2 Running Example

We introduce our approach on the example of a simple web service. The example is a translation from PHP with embedded SQL commands into SmpSL of code excerpts from the Panda web-administrator. The web service provides several services implemented in dedicated functions for subscribing a user to a newsletter, deleting a newsletter, making a user an admin of a newsletter, sending emails to all subscribed users of a newsletter, etc. We illustrate our verification methodology by exposing an error in the Panda web-administrator. The verification methodology we envision in this paper consists of (1) maintaining database invariants and (2) verifying a contract specification for each function of the web service.

The database contains several tables including $NS = \mathit{NewsletterSubscription}$ with attributes *nwl*, *user*, *subscribed* and *code*. The database is a structure whose universe is partitioned into three sets: $\mathbf{dom}^{\mathrm{U}}$, $\mathbf{bool}^{\mathrm{B}}$, and $\mathbf{codes}^{\mathrm{B}}$. The attributes *nwl* and *user* range over the finite set $\mathbf{dom}^{\mathrm{U}}$, the attribute *subscribed* ranges over $\mathbf{bool}^{\mathrm{B}} = \{\mathit{true}, \mathit{false}\}$, and the attribute *code* ranges over the fixed finite set $\mathbf{codes}^{\mathrm{B}}$. The superscripts in $\mathbf{dom}^{\mathrm{U}}$, $\mathbf{bool}^{\mathrm{B}}$, and $\mathbf{codes}^{\mathrm{B}}$ serve to indicate that the domain $\mathbf{dom}^{\mathrm{U}}$ is unbounded, while the Boolean domain and the domain of codes are bounded (i.e. of fixed finite size). When $s = \mathit{true}$, $(n, u, s, c) \in NS$ signifies that the *user* $u$ is *subscribed* to the *newsletter* $n$. The process of being (un)subscribed from/to a newsletter requires an intermediary confirmation step in which the confirm code $c$ plays a role.

Figure 1 provides the functions `subscribe`, `unsubscribe`, and `confirm` translated manually into SmpSL. [1] The comments in quotations `// "..."` originate from the PHP source code. The intended use of these functions is as follows: In order to subscribe a user $u$ to a newsletter $n$, the function `subscribe` is called with inputs $n$ and $u$ (for example by a web interface operated by the newsletter admin or by the user). `subscribe` stores the tuple $(n, u, \mathit{false}, \mathit{new_code})$ in $NS$, where *new_code* is a confirmation code which does not occur in the database, and an email containing a confirmation URL is sent to the user $u$. Visiting the URL triggers a call to `confirm` with input *new_code*, which subscribes $u$ to $n$ by replacing the tuple $(n, u, \mathit{false}, \mathit{new_code})$ of $NS$ to with $(n, u, \mathit{true}, \texttt{nil})$. For `unsubscribe` the process is similar, and crucially, `unsubscribe` uses the same `confirm` function. `confirm` decides whether to subscribe or unsubscribe according to whether $n$ is currently subscribed to $u$. The `CHOOSE` command selects one row non-deterministically.

The database preserves the invariant

$$\mathit{Inv} \quad = \quad \forall_{\mathbf{d}} x, y. \forall_{\mathbf{b}} s_1, s_2. \forall_{\mathbf{c}} c_1, c_2. \left( (s_1 = s_2 \wedge c_1 = c_2) \vee \bigvee\nolimits_{i=1,2} \neg NS(x, y, s_i, c_i) \right) \qquad (1)$$

*Inv* says that the pair $(n, u)$ of newsletter and user is a key of the relation $NS$. The subscripts of the quantifiers denote the domains over which the quantified variables range. In our verification methodology we add invariants as additional conjuncts to the pre- and post-conditions of every function. In this way invariants strengthen the pre-conditions and can be used to prove the post-conditions of the functions. On the other hand, the post-conditions require to re-establish the validity of the invariants.

Figure 2 provides pre- and post-conditions $\mathrm{pre}_{\mathtt{f}}$ and $\mathrm{post}_{\mathtt{f}}$ for each of the three functions `f`. The relation names **d**, **b**, and **c** are interpreted as the sets $\mathbf{dom}^{\mathrm{U}}$, $\mathbf{bool}^{\mathrm{B}}$, and $\mathbf{codes}^{\mathrm{B}}$, respectively. Proving correctness amounts to proving the correctness of each of the *Hoare triples* $\{\mathrm{pre}_{\mathtt{f}} \wedge \mathit{Inv}\}$ `f` $\{\mathrm{post}_{\mathtt{f}} \wedge \mathit{Inv}\}$. Each Hoare triple specifies a contract: after every

[1] The reader may wish to compare the SmpSL implementation of `confirm` to the PHP implementation in PANDA, provided in Appendix A.1.

```
subscribe(n,u):
    A = SELECT * FROM NS WHERE user = u AND nwl = n;
    if (A != empty) exit; // "This address is already registered to this newsletter."
    INSERT (n,u,false,new_code) INTO NS;
    // Send confirmation email to u

unsubscribe(n,u):
    A = SELECT * FROM NS WHERE user = u AND nwl = n;
    if (A = empty) exit; // "This address is not registered to this newsletter."
    UPDATE NS SET code = new_code WHERE user = u AND nwl = n
    // Send confirmation email to u

confirm(cd):
    A = SELECT subscribe FROM NS WHERE code = cd;
    if (A = empty) exit; //"No such code"
    s1 = CHOOSE A;
    if (s1 = false) UPDATE NS SET subscribed = true, code = nil WHERE code = cd
    else DELETE FROM NS WHERE code = cd;
```

**Figure 1** Running Example: SmpSL code.

execution of `f`, the condition $\text{post}_{\texttt{f}} \wedge \mathit{Inv}$ should be satisfied if $\text{pre}_{\texttt{f}} \wedge \mathit{Inv}$ was satisfied before executing `f`. $\text{pre}_{\texttt{subscribe}}$ and $\text{pre}_{\texttt{unsubscribe}}$ express that *new_code* is an unused non-nil code and that $NS_{gh}$ is equal to $NS$. $NS_{gh}$ is a *ghost table*, used in the post-conditions to relate the state before the execution of the function to the state after the execution. $NS_{gh}$ does not occur in the functions and is not modified. $\text{post}_{\texttt{subscribe}}$ and $\text{post}_{\texttt{unsubscribe}}$ express that $NS$ is obtained from $NS_{gh}$ by inserting or updating a row satisfying `user` $=$ `u` `AND` `nwl` $=$ `n` whenever the `exit` command is not executed. The intended behavior of `confirm` depends on which function created *cd*. $\text{pre}_{\texttt{confirm}}$ introduces a Boolean ghost variable $sub_{gh}$ whose value is true (respectively false) if *cd* was generated as a new code in `subscribe` (respectively `unsubscribe`). $sub_{gh}$ does not occur in `confirm`. $\text{post}_{\texttt{confirm}}$ express that, when $sub_{gh}$ is true, $NS$ is obtained from $NS$ by toggling the value of the column *subscribed* from false to true in the $NS_{gh}$ row whose confirm code is *cd*; when $sub_{gh}$ is false, $NS$ is obtained from $NS_{gh}$ by deleting the row with confirm code *cd*.

Let us now describe the error which prevents `confirm` from satisfying its specification. Consider the following scenario. First, `subscribe` is called and then `unsubscribe`, both with the same input $n$ and $u$. Two confirm codes are created: $c_{\texttt{s}}$ by `subscribe` and $c_{\texttt{u}}$ by `unsubscribe`. At this point, $NS$ contains a single row for the newsletter $n$ and user $u$ namely $(n, u, \mathit{false}, c_{\texttt{u}})$. The user receives two confirmation emails containing the codes $c_{\texttt{s}}$ and $c_{\texttt{u}}$. Clicking on the confirmation URL for $c_{\texttt{s}}$ (i.e. running $\texttt{confirm}(c_{\texttt{s}})$) has no effect since $c_{\texttt{s}}$ does not occur in the database. However, clicking on the confirmation URL for $c_{\texttt{s}}$ results in subscribing $u$ to $n$. This is an error, since confirming a code created in `unsubscribe` should not lead to a subscription.

Our tool automatically checks whether the program satisfies its specification. If not, the programmer or verification engineer may try to refine the specification to adhere more closely to the intended behavior (e.g. by adding an invariant). In this case, the program is in fact incorrect, so no meaningful correct specification can be written for it.

In Section 3.3 we describe a *weakest-precondition calculus* $\text{wp}[\![\cdot]\!]$ which allows us to automatically derive the weakest precondition for a post-condition with regard to a SmpSL program. For our example functions `f`, $\text{wp}[\![\cdot]\!]$ allows us to automatically derive $\text{wp}[\![\texttt{f}]\!]\text{post}_{\texttt{f}}$. The basic property of the weakest precondition is that $\text{post}_{\texttt{f}}$ holds after `f` has executed iff $\text{wp}[\![\texttt{f}]\!]\text{post}_{\texttt{f}}$ held immediately at the start of the execution. It then remains to show that the pre-condition $\text{pre}_{\texttt{f}}$ implies $\text{wp}[\![\texttt{f}]\!]\text{post}_{\texttt{f}}$. This amounts to checking the validity of the

$$
\begin{array}{lcl}
\text{pre}_{\texttt{g}} & = & NS = NS_{gh} \wedge \text{good-code}(new_code) \\
\text{good-code}(c') & = & \mathbf{c}(c') \wedge (c' \neq nil) \wedge \forall_{\mathbf{d}} x, y. \forall_{\mathbf{b}} s. \neg NS(x, y, s, c') \\
\text{post}_{\texttt{g}} & = & \forall_{\mathbf{d}} x, y. \forall_{\mathbf{b}} s. \forall_{\mathbf{c}} c. \, NS(x, y, s, c) \leftrightarrow (\varphi_{\texttt{g},1} \vee \varphi_{\texttt{g},2}) \\
\varphi_{\texttt{subscribe},1} & = & NS_{gh}(x, y, s, c) \\
\varphi_{\texttt{subscribe},2} & = & (n = x) \wedge (u = y) \wedge (s = false) \wedge (c = new_code) \\
 & & \wedge \neg \exists_{\mathbf{b}} s'. \exists_{\mathbf{c}} c'. \, NS_{gh}(n, u, s', c') \\
\varphi_{\texttt{unsubscribe},1} & = & (n \neq x) \wedge (u \neq y) \wedge NS_{gh}(x, y, s, c) \\
\varphi_{\texttt{unsubscribe},2} & = & (n = x) \wedge (u = y) \wedge (c = new_code) \wedge \exists_{\mathbf{c}} c'. \, NS_{gh}(n, u, s, c') \\
 & & \\
\text{pre}_{\texttt{confirm}} & = & NS = NS_{gh} \wedge \mathbf{b}(sub_{gh}) \\
\text{post}_{\texttt{confirm}} & = & \bigwedge_{tt \in \mathbf{b}} sub_{gh} = tt \rightarrow (\forall_{\mathbf{d}} x, y. \forall_{\mathbf{b}} s. \forall_{\mathbf{c}} c. \, NS(x, y, s, c) \leftrightarrow \psi_{tt}) \\
\psi_{false} & = & cd \neq c \wedge NS_{gh}(x, y, s, c) \\
\psi_{true} & = & cd \neq c \wedge NS_{gh}(x, y, s, c) \vee (c = nil \wedge s = true \wedge NS_{gh}(x, y, false, cd))
\end{array}
$$

**Figure 2** Running Example: Pre- and post-conditions. `g` is either `subscribe` or `unsubscribe`.

verification conditions $VC_{\texttt{f}} = \text{pre}_{\texttt{f}} \rightarrow \text{wp}[\![\texttt{f}]\!]\text{post}_{\texttt{f}}$.

Our reasoner for $\text{FO}^2$ sentences is the back-end for our verification tool. The specification in this example is all in $\text{FO}^2_{\text{BD}}$. The weakest precondition of a SmpSL program applied to a $\text{FO}^2_{\text{BD}}$ sentence gives again a $\text{FO}^2_{\text{BD}}$sentence. Hence $VC_{\texttt{f}}$ are all in $\text{FO}^2_{\text{BD}}$. Automatically deciding the validity of $\text{FO}^2_{\text{BD}}$ sentences using our $\text{FO}^2$ decision procedure is described in Section 4. Recall that $\mathbf{codes}^{\text{B}}$ is of fixed finite size. Here $|\mathbf{codes}^{\text{B}}| = 3$ is sufficient to detect the error. Observe that the same confirm code may be reused once it is replaced with `nil` in `confirm`, so the size of the database is unbounded. The size of $\mathbf{codes}^{\text{B}}$ must be chosen manually when applying our automatic tool.

A simple way to correct the error in `confirm` is by adding $sub_{gh}$ as a second argument of `confirm` and replacing `if` $(\texttt{s1} = \texttt{false}) \cdots$ with `if` $(\texttt{sub}_{\texttt{gh}} = \texttt{false}) \cdots$. Since $s_1$ is no longer used, the CHOOSE command can be deleted. The value of $sub_{gh}$ received by `confirm` is set correctly by `subscribe` and `unsubscribe`. With these changes, the error is fixed and `confirm` satisfies its specification. In the scenario from above, the call to `confirm` with $c_{\texttt{s}}$ and $sub_{gh} = true$ leaves the database unchanged, while the call to `confirm` with $c_{\texttt{u}}$ and $sub_{gh} = false$ deletes the row $(n, u, false, c_{\texttt{u}})$.

# 3 Verification of SmpSL Programs

Here we introduce our programming language and our verification methodology. We introduce the SQL fragment SmpSQL in Section 3.1 and the scripting language SmpSL in Section 3.2. In Section 3.3 we explain the weakest precondition transformer of SmpSL, and we show how discharging verification conditions of $\text{FO}^2_{\text{BD}}$ specification reduces to reasoning in $\text{FO}^2$.

## 3.1 The SQL fragment SmpSQL

### 3.1.1 Data model of SmpSQL

The data model of SmpSQL is based on the presentation of the relational model in Chapter 3.1 of [2]. We assume finite sets of $\mathbf{dom}^{\text{B}}_1, \ldots, \mathbf{dom}^{\text{B}}_s$ called the *bounded domains* and an infinite set $\mathbf{dom}^{\text{U}}$ called the *unbounded domain*. The domains are disjoint. We assume three disjoint countably infinite sets: the set of attributes $\mathbf{att}$, the set of relation names $\mathbf{relnames}$, and the set of variables $\mathbf{SQLvars}$. We assume a function $\mathbf{sort} : \mathbf{att} \rightarrow \{\mathbf{dom}^{\text{U}}, \mathbf{dom}^{\text{B}}_1, \ldots, \mathbf{dom}^{\text{B}}_s\}$.

A ***table*** or a ***relation schema*** is a relation name and a finite sequence of attributes. The attributes are the names of the columns of the table. The ***arity*** $\mathrm{ar}(R)$ of a relation schema $R$ is the number of its attributes. A ***database schema*** is a non-empty finite set of tables.

A ***database instance*** $\mathcal{I}$ of a database schema $\mathbf{R}$ is a many-sorted structure with finite domains $\mathbf{dom}_0 \subseteq \mathbf{dom}^{\mathrm{U}}$ and $\mathbf{dom}_j = \mathbf{dom}^{\mathrm{B}}_j$ for $1 \leq j \leq s$. We denote by $\mathbf{sort}_{\mathcal{I}}$ the function obtained from $\mathbf{sort}$ by setting $\mathbf{sort}_{\mathcal{I}}(att) = \mathbf{dom}_0$ whenever $\mathbf{sort}(att) = \mathbf{dom}^{\mathrm{U}}$. The relation schema $R = (relname, att_1, \ldots, att_e)$ is interpreted in $\mathcal{I}$ as a relation $R^{\mathcal{I}} \subseteq \mathbf{sort}_{\mathcal{I}}(att_1) \times \cdots \times \mathbf{sort}_{\mathcal{I}}(att_e)$. A ***row*** is a tuple in a relation $R^{\mathcal{I}}$.

A database schema $\mathbf{R}$ is **valid** for SmpSQL if for all relation schemas $R$ with attributes $att_1, \ldots, att_e$ in $\mathbf{R}$, there are at most two attributes $att_j$ for which $\mathbf{sort}(att_j) = \mathbf{dom}^{\mathrm{U}}$. In the sequel we assume that all database schemas are valid. The SmpSQL commands will be allowed to use variables from **SQLvars**. We denote members of **SQLvars** by $p$, $p_1$, etc.

### 3.1.2 Queries in SmpSQL

Given a relation schema $R$ and attributes $att_1, \ldots, att_n$ of $R$, the syntax of SELECT is:

$$
\begin{array}{lcl}
\langle \texttt{Select} \rangle & ::= & \texttt{SELECT att}_{\texttt{a}_1}, \ldots, \texttt{att}_{\texttt{a}_i} \texttt{ FROM R WHERE } \langle \texttt{Condition} \rangle \\
\langle \texttt{Condition} \rangle & ::= & \texttt{att}_{\texttt{b}_1}, \ldots, \texttt{att}_{\texttt{b}_j} \texttt{ IN } \langle \texttt{Select} \rangle \mid \langle \texttt{Condition} \rangle \texttt{ AND } \langle \texttt{Condition} \rangle \mid \\
 & & \langle \texttt{Condition} \rangle \texttt{ OR } \langle \texttt{Condition} \rangle \mid \texttt{NOT } \langle \texttt{Condition} \rangle \mid \texttt{att}_{\texttt{m}} = \texttt{p}
\end{array}
$$

where $p$ is a variable and $1 \leq m, a_1, \ldots, a_i, b_1, \ldots, b_j \leq n$. The semantics of $\langle Select \rangle$ is the set of tuples from the projection of $R$ on $att_{a_1}, \ldots, att_{a_i}$ which satisfy $\langle Condition \rangle$. The condition $\texttt{att}_{\texttt{m}} = \texttt{p}$ indicates that the set of rows of $R$ in which the attribute $att_m$ has value $p$ is selected. The condition $\texttt{att}_{\texttt{b}_1}, \ldots, \texttt{att}_{\texttt{b}_i}$ IN $\langle \texttt{Select} \rangle$ selects the set of rows of $R$ in which $att_{b_1}, \ldots, att_{b_i}$ are mapped to one of the tuples queried in the nested query $\langle Select \rangle$.

### 3.1.3 Data-manipulating commands in SmpSQL

SmpSQL supports the three primitive commands INSERT, UPDATE, and DELETE.

Let $R$ be a relation schema with attributes $att_1, \ldots, att_n$. Let $p, p_1, \ldots, p_n$ be variables from **SQLvars**. The syntax of the primitive commands is:

$$
\begin{array}{lcl}
\langle \texttt{Insert} \rangle & ::= & \texttt{INSERT } (\texttt{p}_1, \ldots, \texttt{p}_n) \texttt{ INTO R} \\
\langle \texttt{Update} \rangle & ::= & \texttt{UPDATE R SET att}_{\texttt{m}} = \texttt{p WHERE } \langle \texttt{Condition} \rangle \\
\langle \texttt{Delete} \rangle & ::= & \texttt{DELETE FROM R WHERE } \langle \texttt{Condition} \rangle
\end{array}
$$

The semantics of INSERT, UPDATE and DELETE is given in the natural way. We allow update commands which set several attributes simultaneously. We assume that the data manipulating commands are used in a domain-correctness fashion, i.e. `INSERT` and `UPDATE` may only assign values from $\mathbf{sort}(att_k)$ to any attribute $att_k$.

## 3.2 The script language SmpSL

### 3.2.1 Data model of SmpSL

The data model of SmpSL extends that of SmpSQL with constant names and additional relation schemas. We assume a countably infinite set of constant names **connames**, which is disjoint from $\mathbf{att}, \mathbf{dom}^{\mathrm{U}}, \mathbf{dom}^{\mathrm{B}}_1, \ldots, \mathbf{dom}^{\mathrm{B}}_s, \mathbf{relnames}$ but contains **SQLvars**.

A ***state schema*** is a database schema $\mathbf{R}$ expanded with a tuple of constant names $\overline{const}$. A ***state*** interprets a state schema. It consists of a database instance $\mathcal{I}$ expanded with a tuple

of universe elements $\overline{const}^{\mathcal{I}}$ interpreting $\overline{const}$. In programs, the constant names play the role of local variables, domain constants (e.g. *true* and *true*) and of inputs to the program[2].

### 3.2.2 SmpSL programs

The syntax of SmpSL is given by

$$
\begin{array}{lcl}
\langle \texttt{Program} \rangle & ::= & \langle \texttt{Command} \rangle \mid \langle \texttt{Program} \rangle \,;\, \langle \texttt{Command} \rangle \\
\langle \texttt{Command} \rangle & ::= & \langle \texttt{Insert} \rangle \mid \langle \texttt{Update} \rangle \mid \langle \texttt{Delete} \rangle \mid \texttt{R} = \langle \texttt{Select} \rangle \mid \bar{\texttt{d}} = \texttt{CHOOSE R} \mid \\
 & & \texttt{if (cond)} \ \langle \texttt{Program} \rangle \ \texttt{else} \ \langle \texttt{Program} \rangle \mid \texttt{if (cond) exit}
\end{array}
$$

Every data-manipulating command $C$ of SmpSQL is a SmpSL command. The semantics of $C$ in SmpSL is the same as in SmpSQL, with the caveat that the variables receive their values from their interpretations (as constant names) in the state, and $C$ is only legal if all the variables of $C$ indeed appear in the state schema as constant names.

The command $\texttt{R} = \langle \texttt{Select} \rangle$ assigns the result of a SmpSQL query to a relation schema $R \in \mathbf{R}$ whose arity and attribute sorts match the select query. Executing the command in a state $(\mathcal{I}, \overline{const}^{\mathcal{I}})$ sets $R^{\mathcal{I}}$ to the relation selected by $S$, leaving the interpretation of all other names unchanged. The variables in the query receive their values from their interpretations in the state, and for the command to be legal, all variables in the query must appear in the state schema as constant names.

Given a relation schema $R \in \mathbf{R}$ with attributes $att_1, \ldots, att_n$ and a tuple $\bar{d} = (d_1, \ldots, d_n)$ of constant names from $\overline{const}$, $\bar{\texttt{d}} = \texttt{CHOOSE R}$ is a SmpSL command. If $R^{\mathcal{I}}$ is empty, the command has no effect. If $R^{\mathcal{I}}$ is not empty, executing this command sets $(d_1^{\mathcal{I}}, \ldots, d_n^{\mathcal{I}})$ to the value of a non-deterministically selected row from $R^{\mathcal{I}}$.

The branching commands have the natural semantics. Two types of branching conditions `cond` are allowed: $(R = \texttt{empty})$ and $(R \mathrel{!}= \texttt{empty})$, which check whether $R^{\mathcal{I}}$ is the empty set, and $(c_1 = c_2)$ and $(c_1 ! = c_2)$, which check whether $c_1^{\mathcal{I}} = c_2^{\mathcal{I}}$.

See Fig. 1 for examples of SmpSL programs.

## 3.3 Verification of SmpSL programs

### 3.3.1 SQL and FO

It is well-established that a core part of SQL is captured by FO by Codd's classical theorem relating the expressive power of relational algebra to relational calculus. While SQL goes beyond FO in several aspects, such as aggregation, grouping, and arithmetic operations (see [27]), these aspects are not allowed in SmpSQL. Hence, FO is especially suited for reasoning about SmpSQL and SmpSL.

The notions of state schema and state fit naturally in the syntax and semantics of FO. In the sequel, a ***vocabulary*** is a tuple of relation names and constant names. For a FO-formula $\psi$, we write $\mathrm{voc}(\psi)$ for the vocabulary consisting of the relation names and constant names in $\psi$. Every state schema $\mathbf{R}$ is a vocabulary. A state $(\mathcal{I}, \overline{const}^{\mathcal{I}})$ interpreting a state schema $\mathbf{R}$ and a tuple of constant names $\overline{const}$ is an $\langle \mathbf{R}, \overline{const} \rangle$-structure.

$$
\begin{array}{lcl}
[\![\texttt{att}_\texttt{i} = \texttt{c}]\!]^R & \widehat{=} & v_i = c \\
[\![\texttt{att}_{\texttt{b}_1}, \ldots, \texttt{att}_{\texttt{b}_j}\ \texttt{IN}\ \texttt{S}_1]\!]^R & \widehat{=} & [\![S_1]\!][v_{b_k}/v_k : 1 \le k \le j] \\
[\![\texttt{cond}_1\ \texttt{AND}\ \texttt{cond}_2]\!]^R & \widehat{=} & [\![cond_1]\!]^R \wedge [\![cond_2]\!]^R \\
[\![\texttt{cond}_1\ \texttt{OR}\ \texttt{cond}_2]\!]^R & \widehat{=} & [\![cond_1]\!]^R \vee [\![cond_2]\!]^R \\
[\![\texttt{NOT}\ \texttt{cond}_1]\!]^R & \widehat{=} & \neg [\![cond_1]\!]^R
\end{array}
$$

$[\![\texttt{SELECT att}_{\texttt{a}_1}, \ldots, \texttt{att}_{\texttt{a}_i}\ \texttt{FROM R WHERE cond}]\!] \widehat{=} (\exists v_{a_{i+1}}, \ldots, v_{a_n} R(\bar{v}) \wedge [\![cond]\!]^R)[v_\ell/v_{a_\ell} : 1 \le \ell \le i]$
where $\{a_1, \ldots, a_n\} = \{1, \ldots, n\}$

$$
\begin{array}{lcl}
\text{wp}[\![\texttt{INSERT}\ (\texttt{c}_1, \ldots, \texttt{c}_\texttt{n})\ \texttt{INTO R}]\!]Q & \widehat{=} & Q\big[R(\bar{\alpha}) \vee \bigwedge_{i=1}^n \alpha_i = c_i \ \big/\ R(\bar{\alpha})\big] \\
\text{wp}[\![\texttt{DELETE FROM R WHERE cond}]\!]Q & \widehat{=} & Q\big[R(\bar{\alpha}) \wedge \neg [\![cond]\!]^R[\alpha_i/v_i : 1 \le i \le n] \ \big/\ R(\bar{\alpha})\big] \\
\text{wp}[\![\texttt{UPDATE R SET att}_\texttt{j} = \texttt{c WHERE cond}]\!]Q & \widehat{=} & Q\big[R(\bar{\alpha}) \wedge \neg [\![cond]\!]^R[\alpha_i/v_i : 1 \le i \le n] \vee \\
 & & \quad \exists v_j R(\overline{\alpha^j}) \wedge [\![cond]\!]^R[\alpha_i^j/v_i : 1 \le i \le n] \wedge \alpha_j = c \\
 & & \quad \big/\ R(\bar{\alpha})\big]
\end{array}
$$

**Figure 3** Rules for weakest precondition for SmpSQL basic commands. We denote by $R$ a relation schema with attributes $\langle att_1, \ldots, att_n \rangle$. We write $\alpha_i^j$ for $\alpha_i$ if $i \neq j$, and for $v_i$ if $i = j$. We denote $\bar{v} = (v_1, \ldots, v_n)$, $\bar{\alpha} = (\alpha_1, \ldots, \alpha_n)$, and $\overline{\alpha^j} = (\alpha_1^j, \ldots, \alpha_n^j)$. Note that each of the last three rows $Q\big[expr(\alpha)\big/R(\bar{\alpha})\big]$ substitutes every occurrence of $R$ with an updated expression $expr$.

### 3.3.2 Hoare verification of SmpSL programs and weakest precondition

Hoare logic is a standard program verification methodology [23]. Let $P$ be a SmpSL program and let $\varphi_{pre}$ and $\varphi_{post}$ be FO-sentences. A ***Hoare triple*** is of the form $\{\varphi_{pre}\}P\{\varphi_{post}\}$. A Hoare triple is a *contract* relating the state before the program is run with the state afterward. The goal of the verification process is to prove that the contract is correct.

Our method of proving that a Hoare triple is valid reduces the problem to that of finite satisfiability of a FO-sentence. We compute the ***weakest precondition*** $\text{wp}[\![P]\!]\varphi_{post}$ of $\varphi_{post}$ with respect to the program $P$. The weakest precondition transformer was introduced in Dijkstra's classic paper [17], c.f. [24]. Let $\mathcal{A}_P$ denote the state after executing $P$ on the initial state $\mathcal{A}$. The main property of the weakest precondition is: $\mathcal{A}_P \models \varphi_{post}$ iff $\mathcal{A} \models \text{wp}[\![P]\!]\varphi_{post}$. Using $\text{wp}[\![\cdot]\!]$ we can rephrase the problem of whether the Hoare triple $\{\varphi_{pre}\}P\{\varphi_{post}\}$ is valid in terms of FO reasoning on finite structures: *Is the FO-sentence* $\varphi_{pre} \to \text{wp}[\![P]\!]\varphi_{post}$ *a tautology?* Equivalently, *is the FO-sentence* $\varphi_{pre} \wedge \neg\text{wp}[\![P]\!]\varphi_{post}$ *unsatisfiable?* Section 3.3.3 discusses the resulting FO reasoning task.

We describe the computation of the weakest precondition inductively for SmpSQL and SmpSL. The weakest precondition for SmpSQL is given in Fig. 3, and for SmpSL in Fig. 4. For SmpSQL conditions, $[\![\cdot]\!]^R$ is a formula with $n$ free first-order variables $v_1, \ldots, v_n$ for a conditional expression in the context of relation schema $R$ of arity $n$. $[\![\texttt{SELECT} \cdots \texttt{FROM R} \cdots]\!]$ is also a formula with free variables $v_1, \ldots, v_n$ describing the rows selected by the SELECT query. The rules $\text{wp}[\![s]\!]Q$ transform a (closed) formula $Q$, which is a postcondition of the command $s$, into a (closed) formula expressing the weakest precondition. The notation $\psi[t/v]$ indicates substitution of all free occurrences of the variable $v$ in $\psi$ by the term $t$.

The notation $\psi\big[\theta(\alpha_1, \ldots, \alpha_n)/R(\alpha_1, \ldots, \alpha_n)\big]$ indicates that any atomic sub-formula of $\psi$ of the form $R(\alpha_1, \ldots, \alpha_n)$ (for any $\alpha_1, \ldots, \alpha_n$) is replaced by $\theta(\alpha_1, \ldots, \alpha_n)$ (with the same $\alpha_1, \ldots, \alpha_n$). The formula $\theta(v_1, \ldots, v_n)$ has $n$ free variables, and $\theta(\alpha_1, \ldots, \alpha_n)$ is obtained by

[2] We deviate from [2] in the treatment of constants in that we do not assume that constant names are always interpreted as *distinct* members of $\mathbf{dom}^{\mathrm{U}}$. This is so since several program variables or inputs can have the same value.

$$
\begin{array}{lcl}
[\![\texttt{c}_1 = \texttt{c}_2]\!] & \widehat{=} & c_1 = c_2 \\
[\![\texttt{c}_1 \texttt{ != } \texttt{c}_2]\!] & \widehat{=} & c_1 \neq c_2 \\
[\![\texttt{R != empty}]\!] & \widehat{=} & \exists v_1, \ldots, v_n R(v_1, \ldots, v_n) \\
[\![\texttt{R= empty}]\!] & \widehat{=} & \neg\exists v_1, \ldots, v_n R(v_1, \ldots, v_n)
\end{array}
$$

$$
\begin{array}{lcl}
\mathrm{wp}[\![\texttt{R} = \texttt{SELECT} \ \cdots]\!]Q & \widehat{=} & Q\big[[\![\texttt{SELECT} \ \cdots]\!](\alpha_1, \ldots, \alpha_n) \big/ R(\alpha_1, \ldots, \alpha_n)\big] \\
\mathrm{wp}[\![(\texttt{d}_1, \ldots, \texttt{d}_n) = \texttt{CHOOSE R}]\!]Q & \widehat{=} & \forall u_1, \ldots, u_n \big(R(u_1, \ldots, u_n) \rightarrow Q[u_i/d_i : 1 \leq i \leq n]\big) \\
\mathrm{wp}[\![\texttt{if cond } \texttt{s}_1 \texttt{ else } \texttt{s}_2]\!]Q & \widehat{=} & (\neg[\![cond]\!] \wedge \mathrm{wp}[\![s_2]\!]Q) \vee ([\![cond]\!] \wedge \mathrm{wp}[\![s_1]\!]Q)
\end{array}
$$

**Figure 4** Rules for weakest precondition construction for SmpSL basic commands. The weakest precondition of $\texttt{if cond exit}; \texttt{s}_2$ is the same as that of $\texttt{if !cond } \texttt{s}_2$.

substituting each $v_i$ into $\alpha_i$. The $\alpha_i$ may be variables or constant names.

The weakest precondition of a SmpSL program is obtained by applying the weakest precondition of its commands.

### 3.3.3 The specification logic $\mathrm{FO}^2_{\mathrm{BD}}$ and decidability of verification

As discussed in Section 3.3.2, using the weakest precondition, the problem of verifying Hoare triples can be reduced to the problem of checking satisfiability of a FO-sentence by a finite structure. While this problem is not decidable in general by Trakhtenbrot's theorem, it is decidable for a fragment of FO we denote $\mathrm{FO}^2_{\mathrm{BD}}$, which extends the classical *two-variable fragment* $\mathrm{FO}^2$. The logic $\mathrm{FO}^2$ is the set of all FO formulas which use only variables the variables $x$ and $y$. The vocabularies of $\mathrm{FO}^2$-sentences are not allowed function names, only relation and constanst names. Note $\mathrm{FO}^2$ cannot express that a relation name is interpreted as a function. $\mathrm{FO}^2$ contains the equality symbol $=$. $\mathrm{FO}^2_{\mathrm{BD}}$ extends $\mathrm{FO}^2$ by allowing quantification on an unbounded number of variables, under the restriction that all variables besides from $x$ and $y$ range over the bounded domains only.

$\mathrm{FO}^2_{\mathrm{BD}}$ is the language of our invariants and pre- and postconditions, see Eq. (1) and Fig. 2 in Section 2. An important property of $\mathrm{FO}^2_{\mathrm{BD}}$ is that it is essentially closed under taking weakest precondition according to Figs. 3 and 4 since all relation schemas in a (valid) database schema have at most 2 attributes whose sort is $\mathbf{dom}^{\mathrm{U}}$. We reduce the task of reasoning over $\mathrm{FO}^2_{\mathrm{BD}}$ to reasoning over $\mathrm{FO}^2$.

▶ **Theorem 1.** *Let $\{\varphi_{pre}\}P\{\varphi_{post}\}$ be a Hoare triple such that both $\varphi_{pre}$ and $\varphi_{post}$ belong to $FO^2_{\mathrm{BD}}$. The problem of deciding whether $\{\varphi_{pre}\}P\{\varphi_{post}\}$ is valid is decidable.*

**(sketch).** By Section 3.3.2, $\{\varphi_{pre}\}P\{\varphi_{post}\}$ is valid iff $\theta = \neg(\varphi_{pre} \wedge \neg \mathrm{wp}[\![P]\!]\varphi_{post})$ is satisfiable by a finite structure. We take the simplifying assumption that in all the tables, the sort of the first and second attributes $att_1$ and $att_2$ is $\mathbf{dom}^{\mathrm{U}}$. This assumption does not effect the expressive power of SmpSL. Examination of the weakest precondition rules in Figs. 3 and 4 reveals that the only variables ranging over the unbounded domain are $v_1$ and $v_2$. Let $\theta'$ be the $\mathrm{FO}^2_{\mathrm{BD}}$ sentence obtained from $\theta$ by substituting $v_1$ and $v_2$ with $x$ and $y$ respectively, and restricting the range of the quantifiers appropriately: for a command manipulating or querying a table $R$ with attributes $att_1, \ldots, att_n$ in Figs. 3 and 4, each quantifier $\forall v_k$ or $\exists v_k$ is replaced with $\forall_{\mathbf{sort}(att_k)} v_k$ or $\exists_{\mathbf{sort}(att_k)} v_k$. We compute an $\mathrm{FO}^2$ sentence $\theta''$ which is equivalent to $\theta'$ by hard-coding the bounded domains. Every table $T$ which contains an attribute $att$ with $\mathbf{sort}(att) = \mathbf{dom}^{\mathrm{B}}_j$ of size $d$ is replaced with $d$ tables $T_1, \ldots, T_d$ which do not have the attribute $att$. This change is reflected in $\theta''$, e.g. existential quantification is replaced with disjunction. By the decidability of finite satisfiability of $\mathrm{FO}^2$-sentences, we get that the problem of deciding whether $\{\varphi_{pre}\}P\{\varphi_{post}\}$ is valid is decidable. ◀

# 4 FO$^2$ Reasoning

## 4.1 The bounded model property of FO$^2$

Section 4 is devoted to our algorithm for FO$^2$ finite satisfiability. The main ingredient for this algorithm is the *bounded model property*, which guarantees that if an FO$^2(\tau)$ sentence $\phi$ over vocabulary $\tau$ is satisfiable by any $\tau$-structure – finite or infinite – it is satisfiable by a finite $\tau$-structure whose cardinality is bounded by a computable function of $\phi$. The bound guaranteed in the first decidability proof of the finite satisfiability problem by Mortimer [29] was doubly exponential in the size of the formula. Later, Grädel, Kolaitis and Vardi [21] proved the *exponential model property*, from which we get that the problem is NEXPTIME-complete. The naive NEXPTIME algorithm arising from the exponential model property amounts to computing the exponential bound $bnd(\phi)$ from [21], non-deterministically guessing $t \leq bnd(\phi)$ and a $\tau$-structure $\mathcal{A}$ with universe $\{1, \ldots, t\}$, checking whether $\mathcal{A}$ satisfies $\phi$, and answering accordingly. Since the truth-value of FO-sentences is invariant to $\tau$-isomorphisms, a $\tau$-structure of cardinality at most $bnd(\phi)$ satisfies $\phi$ iff such a structure with universe $\{1, \ldots, t\}$, $t \leq bnd(\phi)$, satisfies $\phi$. Appendix D.3 discusses a more refined version of the bound from [21].

## 4.2 Finite satisfiability using a SAT solver

Our algorithm for FO$^2$ finite satisfiability reduces the problem of finding a satisfying model of cardinality bounded by *bnd* to the satisfiability of a propositional Boolean formula in Conjunctive Normal Form CNF, which is then solved using a SAT solver. The bound in [21] is given for formulas in Scott Normal Form (SNF) only. We use a refinement of SNF we call *Skolemized Scott Normal Form (SSNF)*. The CNF formula we generate encodes the semantics of the sentence $\psi$ on a structure whose universe cardinality is bounded by *bnd*. An early precursor for the use of a SAT solver for finite satsifiability is [28].

### 4.2.1 Skolemized Scott Normal Form

An FO$^2$-sentence is in *Skolemized Scott Normal Form* if it is of the form

$$\forall x \forall y \left( \alpha(x,y) \wedge \bigwedge_{i=1}^{m} F_i(x,y) \rightarrow \beta_i(x,y) \right) \wedge \bigwedge_{i=1}^{m} \forall x \exists y \, F_i(x,y) \tag{2}$$

where $\alpha$ and $\beta_i$, $i = 1, \ldots, m$, are quantifier-free formulas which do not contain any $F_j$, $j = 1, \ldots, m$. Note that $F_i$ are relation names.

▶ Proposition 1. Let $\tau$ be a vocabulary and $\phi$ be a FO$^2(\tau)$-sentence. There are polynomial-time computable vocabulary $\sigma \supseteq \tau$ and FO$^2(\sigma)$-sentence $\psi$ such that (a) $\psi$ is in SSNF; (b) The set of cardinalities of the models of $\phi$ is equal to the corresponding set for $\psi$; and (c) The size of $\psi$ is linear in the size of $\phi$.

Proposition 1 follows from the discussion before Proposition 3.1 in [21], by applying an additional normalization step converting SNF sentences to SSNF sentences.[3][4]

[3] The word Skolemized is used in reference to the standard *Skolemization* process of eliminating existential quantifiers by introducing fresh function names called *Skolem functions*. In our case, since function names are not allowed in our fragment, we introduce the relation names $F_i$, to which we refer as *Skolem relations*. Moreover, we cannot eliminate the existential quantifiers entirely, but only simplify the formulas in their scope to the atoms $F_i(x,y)$.

[4] The linear size of $\psi$ uses our relation symbols have arity at most 2 to get rid of a log factor in [21].

### 4.2.2 The CNF formula

Given the sentence $\psi$ in SSNF from Eq. (2) and a bound $bnd(\psi)$, we build a CNF propositional Boolean formula $C_\psi$ which is satisfiable iff $\psi$ is satisfiable. The formula $C_\psi$ will serve as the input to the SAT solver. First we construct a related CNF formula $B_\psi$. The crucial property of $B_\psi$ is that it is satisfiable iff $\psi$ is satisfiable by a model *of cardinality exactly* $bnd(\psi)$.

It is convenient to assume $\psi$ does not contain constants. If $\psi$ did contain constants $c$, they could be replaced by unary relations $U_c$ of size 1. Being an unary relation of size 1 is definable in $\mathrm{FO}^2$. Any atom containing $c$ cannot use both $x$ and $y$, and hence the universe member interpreting $c$ can be quantified: e.g. $R(x,c)$ is replaced with $\exists y\, U_c(y) \wedge R(x,y)$. Let $\mathrm{const}(\psi)$ be the set of unary relations $U_c$ corresponding to constants.

We start by introducing the variables and clauses which guarantee that $B_\psi$ encodes a structure with the universe $\{1,\dots,bnd(\psi)\}$. Later, we will add clauses to guarantee that this structure satisfies $\psi$. For every unary relation name $U$ in $\psi$ and $\ell_1 \in \{1,\dots,bnd(\psi)\}$, let $v_{U,\ell_1}$ be a propositional variable. For every binary relation name $R$ in $\psi$ and $\ell_1,\ell_2 \in \{1,\dots,bnd(\psi)\}$, let $v_{R,\ell_1,\ell_2}$ be a propositional variable. The variables $v_{U,\ell_1}$ and $v_{R,\ell_1,\ell_2}$ encode the interpretations of the unary and binary relation names $U$ and $R$ in the straightforward way (defined precisely below). Let $V_\psi$ be the set of all variables $v_{U,\ell_1}$ and $v_{R,\ell_1,\ell_2}$.

Given an assignment S to the variables of $V_\psi$ we define the unique structure $\mathcal{A}_\mathrm{S}$ as follows:

1. The universe $A_\mathrm{S}$ of $\mathcal{A}_\mathrm{S}$ is $\{1,\dots,bnd(\psi)\}$;
2. An unary relation name $U$ is interpreted as the set $\{\ell_1 \in A_\mathrm{S} \mid \mathrm{S}(v_{U,\ell_1}) = True\}$;
3. A binary relation name $R$ is interpreted as the set $\{(\ell_1,\ell_2) \in A_\mathrm{S}^2 \mid \mathrm{S}(v_{R,\ell_1,\ell_2}) = True\}$;

For every structure $\mathcal{A}$ with universe $\{1,\dots,bnd(\psi)\}$, there is S such that $\mathcal{A} = \mathcal{A}_\mathrm{S}$.

Before defining $B_\psi$ precisely we can already state the crucial property of $B_\psi$:

▶ Proposition 2. $\psi$ is satisfiable by a structure with universe $\{1,\dots,bnd(\psi)\}$ iff $B_\psi$ is satisfiable.

The formula $B_\psi$ is the conjunction of $B^{\mathrm{eq}}$, $B^{\forall\exists}$, and $B^{\forall\forall}$, described in the following.

**The equality symbol.** The equality symbol requires special attention. Let

$$B^{\mathrm{eq}} = \bigwedge_{1\le \ell_1 \ne \ell_2 \le m} (\neg v_{=,\ell_1,\ell_2}) \wedge \bigwedge_{1\le \ell \le m} v_{=,\ell,\ell}$$

$B^{\mathrm{eq}}$ enforces that the equality symbol is interpreted correctly as the equality relation on universe elements.

**The $\forall\exists$-conjuncts.** For every conjunct $\forall x \exists y\, F_i(x,y)$ and $1 \le \ell_1 \le bnd(\psi)$, let $B^{\forall\exists}_{i,\ell_1}$ be the clause $\bigvee_{\ell_2=1}^{bnd(\psi)} v_{F_i,\ell_1,\ell_2}$. This clause says that there is at least one universe element $\ell_2$ such that $\mathcal{A}_\mathrm{S} \models F(\ell_1,\ell_2)$. Let

$$B^{\forall\exists} = \bigwedge_{1\le i\le m} \bigwedge_{1\le \ell_1 \le bnd(\psi)} B^{\forall\exists}_{i,\ell_1}$$

For every truth-value assignment S to $V_\psi$, $\mathcal{A}_\mathrm{S}$ satisfies $\bigwedge_{i=1}^m \forall x \exists y\, F_i(x,y)$ iff S satisfies $B^{\forall\exists}$.

**The $\forall\forall$-conjunct.** Let $\forall x \forall y\, \alpha'$ be the unique $\forall\forall$-conjunct of $\psi$. For every $1 \le \ell_1,\ell_2 \le bnd(\psi)$, let $\alpha''_{\ell_1,\ell_2}$ denote the propositional formula obtained from the quantifier-free $\mathrm{FO}^2$ formula $\alpha'$ by substituting every atom $a$ with the corresponding propositional variable for $\ell_1$ and $\ell_2$ as follows:

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| $U(x)$ | $\mapsto$ | $v_{U,\ell_1}$, | $R(y,y)$ | $\mapsto$ | $v_{R,\ell_2,\ell_2}$, | $R(x,x)$ | $\mapsto$ | $v_{R,\ell_1,\ell_1}$ |
| $U(y)$ | $\mapsto$ | $v_{U,\ell_2}$, | $R(x,y)$ | $\mapsto$ | $v_{R,\ell_1,\ell_2}$, | $R(y,x)$ | $\mapsto$ | $v_{R,\ell_2,\ell_1}$ |

Let $B^{\forall\forall}_{\ell_1,\ell_2}$ be the Tseitin transformation of $\alpha''_{\ell_1,\ell_2}$ to CNF [35], see also [6, Chapter 2]. The Tseitin transformation introduces a linear number of new variables of the form $u^{\gamma}_{\ell_1,\ell_2}$, one for

| | FO$^2$-solver | Z3 |
|---|---|---|
| web-subscribe | 2.62s | 0.08s |
| web-unsubscribe | 0.779s | OM |
| firewall | 0.876s | OM |
| conf-bid | 0.451s | 0.015s |
| conf-assign | 0.369s | 0.013s |
| conf-display | 0.992s | 0.016s |

incorrect

| | FO$^2$-solver | Z3 |
|---|---|---|
| web-subscribe | 1.07s | 0.1s |
| web-unsubscribe | 8.209s | 0.1s |
| firewall | 2.82s | 0.103s |
| conf-bid | TO | 0.22s |
| conf-assign | 1.196s | 0.2s |
| conf-display | TO | 0.16s |

correct

**Table 1** Running time comparison for example benchmarks

each sub-formula $\gamma$ of $\alpha''_{\ell_1,\ell_2}$. The transformation guarantees that, for every assignment S of $V_\psi$, S satisfies $\alpha''_{\ell_1,\ell_2}$ iff S can be expanded to satisfy $B^{\forall\forall}_{\ell_1,\ell_2}$. Let

$$B^{\forall\forall} = \bigwedge_{1\leq \ell_1,\ell_2 \leq bnd(\psi)} B^{\forall\forall}_{\ell_1,\ell_2}(\ell_1,\ell_2)$$

Appendix D.1 gives the construction of the CNF formula $B^{\forall\forall}$ according to the Tseitin transformation explicitly.

The construction of $B_\psi$ is finished and Proposition 2 holds. Note that [21] guarantees only that $bnd(\psi)$ is an *upper bound* on the cardinality of a satisfying model. Therefore, we build a formula $C_\psi$ based on $B_\psi$ such that $C_\psi$ is satisfiable iff $\psi$ is satisfiable by a structure of cardinality *at most* $bnd(\psi)$. We leave the technical details of the construction of $C_\psi$ to the appendix. The algorithm for finite satisfiability of a FO$^2$-sentence $\phi$ consists of computing the SSNF $\psi$ of $\phi$ and returning the result of a satisfiability check using a SAT solver on $C_\psi$. Both the number of variables and the number of clauses in $C_{\mathrm{Uni}(\psi)}$ are quadratic in $bnd(\psi)$.

# 5 Experimental Results

## 5.1 Details of our tools

The verification condition generator described in Section 3.3.2 is implemented in Java, JFlex and CUP. It is employed to parse the schema, precondition and postcondition and the SmpSL programs. The tool checks that the pre and post conditions are specified in FO$^2$ and that the scheme is well defined. The SMT-LIB v2 [4] standard language is used as the output format of the verification condition generator. We compare the behavior of our FO$^2$-solver with Z3 on the verification condition generator output. The validity of the verification condition can be checked by providing its negation to the SAT solver. If the SAT solver exhibits a satisfying assignment then that serves as counterexample for the correctness of the program. If no satisfying assignment exists, then the generated verification condition is valid, and therefore the program satisfies the assertions. The FO$^2$-solver described in Section 4 is implemented in python and uses pyparsing to parse the SMT-LIB v2 [4] file. The FO$^2$-solver assumes a FO$^2$-sentence as input and uses *Lingeling* [5] SAT solver as a base Solver.

## 5.2 Example applications

We tried our approach with a few programs inspired by real-life applications. The first case study is a simplified version of the newsletter functionality included in the PANDA web

administrator, that was already discussed and is shown in Fig. 1.[5] The second is an excerpt from a firewall that updates a table of which device is allowed to send packets to which other device. See Appendix C for the code and specifications of the firewall. The third is a conference management system with a database of papers, and transactions to manage the review process: reviewers first *bid* on papers from the pool of submissions, with a policy that a users cannot bid for papers with which they are conflicted. The chair then *assigns* reviewers to papers by selecting a subset of the bids. At any time, users can ask to *display* the list of papers, with some details, but the system may hide some confidential information, in particular, users should not be able to see the status of papers before the program is made public. We show how our system detects an information flow bug in which the user might learn that some papers were accepted prematurely by examining the session assignments. This bug is based on a bug we observed in a real system. See Appendix B for the code and specifications of the conference management system. Each example comes with two specifications, one correct and the other incorrect.

The running time in seconds for all of our examples is reported in Table 1. Timeout is set to 60 minutes and denoted as TO. If the solver reaches *out of memory* we mark it as OM. On the set of correct examples, both solvers answer within a few seconds, Z3 terminates within milliseconds, while $FO^2$-solver takes a few seconds and times out on some of them. On the set of incorrect examples, Z3 fails to answer while our solver performs well. Note that correct examples correspond to unsatisfiable $FO^2$-sentences, while incorrect examples correspond to satisfiable $FO^2$-sentences.

## 5.3 Examining scalability

**Inflated examples.** In order to evaluate scalability to large examples we inflated our base examples. For instance, while the *subscribe* example from Table 1 consisted of the subscription of one new email to a mailing-list, Table 2 presents analogous examples in which multiple emails are subscribed to multiple mailing-lists. The column *multiplier* details the number of individual subscriptions in each example program. The *unsubscribe* and *firewall* example programs are inflated similarly (see Appendix C.1).

We have tested both our $FO^2$-solver and Z3 on large examples and the results reported in Table 2. The high-level of the results is similar to the case of the small examples. On the incorrect examples set Z3 continues to fail mostly due to running out of memory, though it succeeds on the subscribe example. On the correct examples set Z3 continues to outperform the $FO^2$-solver.

**Artificial examples.** In addition, we constructed a set of artificial benchmarks comprising of several families of $FO^2$-sentences. Each family is parameterized by a number that controls the size of the sentences (roughly corresponding to the number of quantifiers in the sentence). These problems are inspired by combinatorial problems such as graph coloring and paths. We ran experiments using the $FO^2$-solver and three publicly available solvers: Z3, CVC4 (which are SMT solvers), and Nitpick (a model checker). The results are collected in Table 3. The artificial benchmarks are available at `http://forsyte.at/wp-content/uploads/artificial-smt2.tar.gz`.

[5] We omit the confirmation step due to a missing feature in the implementation of the weakest precondition, however the final version of the tool will support the code from Table 1.

| | | FO$^2$-solver | | | Z3 | | |
|---|---|---|---|---|---|---|---|
| | multiplier | 1 | 10 | 100 | 1 | 10 | 100 |
| incorrect | subscribe | 2.62s | 0.973s | 4.04s | 0.08s | 0.126s | 0.203s |
| | unsubscribe | 0.779s | 0.529s | 1.27s | OM | OM | OM |
| | firewall | 0.876s | 0.723s | 2.251s | OM | OM | OM |
| correct | subscribe | 1.07s | 98.249s | TO | 0.1s | 0.11s | 0.116s |
| | unsubscribe | 8.209s | 456.308s | TO | 0.1s | 0.157s | 0.201s |
| | firewall | 2.82s | 50.142s | 2951.882s | 0.103s | 0.121s | 0.143s |

**Table 2** Running time comparison on inflated examples

| | size | status | Z3 | CVC4 | Nitpick | FO$^2$-solver |
|---|---|---|---|---|---|---|
| 2col | 3 | unsat | 0m0.037s | 0m0.076s | TO | TO |
| | 4 | sat | TO | TO | 0m7.038s | 0m5.433s |
| | 5 | unsat | 0m0.702s | 0m0.477s | TO | TO |
| | 6 | sat | TO | TO | 0m8.973s | 0m9.323s |
| | 10 | sat | TO | TO | 0m37.944s | 0m19.580s |
| | 11 | unsat | 1m32.664s | 0m30.912s | TO | TO |
| | 14 | sat | TO | TO | 2m13.661s | TO |
| | 40 | sat | TO | TO | TO | TO |
| alternating-paths | 2 | sat | 0m0.049s | TO | 0m11.144s | 0m1.105s |
| | 100 | sat | TO | TO | TO | 0m9.671s |
| alternating-simple-paths | 3 | sat | TO | TO | TO | 0m6.754s |
| | 4 | sat | TO | TO | TO | 0m10.128s |
| | 7 | sat | TO | TO | TO | TO |
| | 10 | sat | TO | TO | TO | TO |
| exponential | 3 | sat | TO | TO | 0m12.255s | 0m1.847s |
| | 4 | sat | TO | TO | 0m15.358s | 11m6.482s |
| one-var-alternating-sat | 300 | sat | 0m0.037s | 0m0.497s | 0m11.605s | 0m9.720s |
| one-var-alternating-unsat | 5 | unsat | 0m0.026s | 0m0.073s | 0m22.537s | 0m54.198s |
| one-var-nested-exists-sat | 300 | sat | 0m0.031s | 0m0.045s | 0m7.132s | 0m0.562s |
| one-var-nested-forall-sat | 500 | sat | 0m0.033s | TO | 0m7.183s | 0m7.318s |
| path-unsat | 2 | unsat | 0m0.033s | 0m0.044s | TO | 1m37.099s |
| | 3 | unsat | 0m0.030s | 0m0.062s | TO | 1m35.451s |
| | 6 | unsat | 0m0.037s | 0m0.891s | TO | 1m39.209s |

**Table 3** Running time comparison on artificial benchmarks

**Scalability of FO$^2$-solver.** We shall conclude that the FO$^2$-solver, despite being a proof of concept in python with minimal optimizations, handles well incorrect specifications (satisfiable sentences) and also scales well on them. However it struggles on the correct specifications and does not scale well. This suggests that in future work we may choose to run both our solver and Z3 in parallel and answer according the first answer obtained. We also intend to explore how to improve the performance of our solver in the case of incorrect examples. By construction, whenever FO$^2$-solver finds a satisfying model, its size is at most 4 times that of the minimal model. (The constant 4 can be decreased or increased. )

## 6 Discussion

**Related work.** Verification of database-centric software systems has received increasing attention in recent years [15]. Tools from program analysis and model-checking are used to reason about the correctness of programs which access a database. Unlike our approach, the services accessing the database are usually provided *a priori* in terms of a specification in the

style of a local contract [30, 26]. The code of the services themselves may be automatically synthesized from the specification, cf. e.g. [19, 20, 15, 16]. The focus of verification then is on global temporal properties of the system assuming the local contracts. In contrast, out goal is to verify that the input code (written by a programmer rather than generated automatically) is correct with respect to a local specification. We discuss this also in Section 6.

Several papers use variations of $\mathrm{FO}^2$ to study verification of programs that manipulate relational information. [8] presents a verification methodology based on $\mathrm{FO}^2$, a description logic and a separation logic for analyzing the shapes and content of in-memory data structures. [32] develops a logic similar to $\mathrm{FO}^2$ to reason about shapes of data structures. In both [8] and [32], the focus is on analysis of shapes in dynamically-allocated memory, and databases are not studied. Furthermore, no tools based on these works are available. A description logic related to $\mathrm{FO}^2$ was used in [9] to verify that graph databases preserve the satisfaction of constraints as they evolve. The focus of this work is on the correctness of the database, rather than the programs manipulating it. The verification method suggested was not implemented. In fact, to our knowledge no description logic solver implements reasoning tasks for the description logic counterpart of $\mathrm{FO}^2$ studied in [9], not even solvers for expressive description logics such as SROIQ.

Verification of script programs with embedded queries has revolved around security, see [18]. However, it seems no other work has been done on such programs.

**Conclusion and future work.** We developed a verification methodology for script programs with access to a relational database via SQL. We isolated a simple but useful fragment SmpSQL of SQL and developed a simple script programming language SmpSL on top of it. We have shown that verifying the correctness of SmpSL programs with respect to specifications in $\mathrm{FO}^2_{\mathrm{BD}}$ is decidable. We implemented a solver for the $\mathrm{FO}^2$ finite satisfiability problem, and, based on it, a verification tool for SmpSL programs. Our experimental results are very promising and suggest that our approach has great potential to evolve into a mainstream method for the verification of script programs with embedded SQL statements.

While we believe that many of the SQL statements that appear in real-life programs fall into our fragment SmpSQL it is evident that future tools need to consider all of database usage in real-world programs. In future work, we will explore the extension of SmpSL and SmpSQL. Our next goal is to be able to verify large, real-life script programs such as Moodle [1], whose programming language and SQL statements use e.g. some arithmetic or simple inner joins. To do so, we will adapt our approach from the custom-made syntax of SmpSL to a fragment of PHP. We will both explore decidable logics extending $\mathrm{FO}^2_{\mathrm{BD}}$, and investigate verification techniques based on undecidable logics including the use of first-order theorem provers such as Vampire [33, 25] and abstraction techniques which guarantee soundness but may result in spurious errors [12]. For dealing with queries with transitive closure, it is natural to consider fragments of Datalog [10].

A natural extension is to consider global temporal specifications in addition to local contracts. Here the goal is to verify properties of the system which can be expressed in a temporal logic such as Linear Temporal Logic LTL [31, 11]. The approach surveyed in [15], which explore global temporal specifications of services given in terms of local contracts, may be a good basis for studying global temporal specifications in our context.

Another research direction which emerges from the experiments in Section 5 is to explore how to improve the performance of our $\mathrm{FO}^2$ solver on unsatisfiable inputs.

## References


1 Moodle. http://sourceforge.net/projects/moodle/, 2001–2015.
2 Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of databases*, volume 8. Addison-Wesley Reading, 1995.
3 Clark Barrett, Christopher L. Conway, Morgan Deters, Liana Hadarean, Dejan Jovanović, Tim King, Andrew Reynolds, and Cesare Tinelli. Cvc4. In *Proceedings of the 23rd International Conference on Computer Aided Verification*, CAV'11, pages 171–177, Berlin, Heidelberg, 2011. Springer-Verlag.
4 Clark Barrett, Aaron Stump, and Cesare Tinelli. C.: The SMT-LIB standard: Version 2.0. Technical report, 2010.
5 Armin Biere. Lingeling, Plingeling, PicoSAT and PrecoSAT at SAT race 2010. 2010.
6 Armin Biere, Marijn Heule, and Hans van Maaren. *Handbook of satisfiability*, volume 185. IOS press, 2009.
7 Jasmin Christian Blanchette and Tobias Nipkow. Nitpick: A counterexample generator for higher-order logic based on a relational model finder. In *International Conference on Interactive Theorem Proving*, pages 131–146. Springer, 2010.
8 Diego Calvanese, Tomer Kotek, Mantas Šimkus, Helmut Veith, and Florian Zuleger. Shape and content. In *Integrated Formal Methods*, pages 3–17. Springer, 2014.
9 Diego Calvanese, Magdalena Ortiz, and Mantas Simkus. Evolving graph databases under description logic constraints. In *Description Logics*, pages 120–131, 2013.
10 Stefano Ceri, Georg Gottlob, and Letizia Tanca. What you always wanted to know about datalog (and never dared to ask). *Knowledge and Data Engineering, IEEE Transactions on*, 1(1):146–166, 1989.
11 Edmund M. Clarke, E. Allen Emerson, and A. Prasad Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 8(2):244–263, 1986.
12 Edmund M. Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Counterexample-guided abstraction refinement. In *Computer Aided Verification, 12th International Conference, 2000, Chicago, IL, USA, July 15-19, 2000, Proceedings*, pages 154–169, 2000.
13 Edgar F Codd. *Relational completeness of data base sublanguages*. IBM Corporation, 1972.
14 Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340. Springer, 2008.
15 Alin Deutsch, Richard Hull, and Victor Vianu. Automatic verification of database-centric systems. *SIGMOD Record*, 43(3):5–17, 2014.
16 Alin Deutsch, Monica Marcus, Liying Sui, Victor Vianu, and Dayou Zhou. A verifier for interactive, data-driven web applications. In *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, SIGMOD '05, pages 539–550, New York, NY, USA, 2005. ACM.
17 Edsger W Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Communications of the ACM*, 18(8):453–457, 1975.
18 Michael Felderer, Matthias Büchler, Martin Johns, Achim D Brucker, Ruth Breu, and Alexander Pretschner. Security testing: A survey. *Advances in Computers*, 2015.
19 Mary F. Fernández, Daniela Florescu, Alon Y. Levy, and Dan Suciu. Declarative specification of web sites with Strudel. *VLDB J.*, 9(1):38–55, 2000.
20 Daniela Florescu, Valerie Issarny, Patrick Valduriez, and Khaled Yagoub. Weave: A data-intensive web site management system. In *In Proc. of the Conf. on Extending Database Technology (EDBT*, 2000.
21 Erich Grädel, Phokion G Kolaitis, and Moshe Y Vardi. On the decision problem for two-variable first-order logic. *Bulletin of symbolic logic*, 3(01):53–69, 1997.

**22** Alessandro Grassi and Marco Nenciarini. Panda - the php-based email administrator. `http://panda-admin.sourceforge.net/index.php?mode=home`, 2007–2015.

**23** Charles A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, 1969.

**24** Ranjit Jhala and Rupak Majumdar. Software model checking. *ACM Computing Surveys (CSUR)*, 41(4):21, 2009.

**25** Laura Kovács and Andrei Voronkov. First-order theorem proving and Vampire. In *Computer Aided Verification*, pages 1–35. Springer, 2013.

**26** S. Kumaran, P. Nandi, T. Heath, K. Bhaskaran, and R. Das. Adoc-oriented programming. In *In Symp. on Applications and the Internet (SAINT)*, 2003.

**27** Leonid Libkin. Expressive power of SQL. *Theoretical Computer Science*, 296(3):379–404, 2003.

**28** William Mccune. A davis-putnam program and its application to finite first-order model search: Quasigroup existence problems. Technical report, Argonne National Laboratory, 1994.

**29** Michael Mortimer. On languages with two variables. *Mathematical Logic Quarterly*, 21(1):135–140, 1975.

**30** A. Nigam and N. S. Caswell. Business artifacts: An approach to operational specification. *IBM Systems Journal*, 42, 2003.

**31** Amir Pnueli. The temporal logic of programs. In *Foundations of Computer Science, 1977., 18th Annual Symposium on*, pages 46–57. IEEE, 1977.

**32** Arend Rensink. Canonical graph shapes. In David Schmidt, editor, *Programming Languages and Systems*, volume 2986 of *Lecture Notes in Computer Science*, pages 401–415. Springer Berlin Heidelberg, 2004.

**33** Alexandre Riazanov and Andrei Voronkov. The design and implementation of Vampire. *AI communications*, 15(2, 3):91–110, 2002.

**34** Boris Trakhtenbrot. The impossibility of an algorithm for the decidability problem on finite classes. In *Proceedings of the USSR Academy of Sciences*, volume 70, pages 569–572, 1950.

**35** Grigori S Tseitin. On the complexity of derivation in propositional calculus. In *Automation of reasoning*, pages 466–483. Springer, 1983.

**36** Hugh E Williams and David Lane. *Web database applications with PHP and MySQL*. O'Reilly Media, Inc., 2004.

# A PANDA web administrator

## A.1 PANDA Source Code: confirm.php

We present the code of `newsletters/confirm.php` from PANDA Web Administrator version 1.0rc2 in Fig. 5. The code was translated manually to the confirm function in Fig. 1. The biggest difference between the PHP code and the SmpSL code is that the PHP code uses `dbh->getRow` to perform an SQL query which returns one row, whereas in SmpSL this is divided into two steps: first a SELECT query is executed and then CHOOSE selects one row. Additionally, the PHP code performs some more sanity checks,

```
<?php
// No code supplied? Why are you calling us?
if(!$_GET['code']){
        print "Code not specified."; die;
}

// Include common stuff
require_once('common.php');

// Fetch code from command line
$vcode = $_GET['code'];

// Check if code exists in DB
$row = $dbh->getRow('SELECT * FROM newsletter_addresses
                     WHERE confirm_code = ?', array(md5($vcode)),
                     DB_FETCHMODE_ASSOC);

// If it doesn't exist, die
if(!$row['confirm_code']){
        print "No such code"; die;
}

// If user is not subscribed,
// then the confirmation is to subscribe him, so do it
if($row['subscribed'] == 'f'){
        if(!$dbh->query('UPDATE newsletter_addresses
                         SET subscribed = TRUE, confirm_code = NULL
                         WHERE confirm_code = ?',array(md5($vcode)))){
                print "Error while accessing database,
                       contact system administrator."; die;
        };
} else { // Else, code is to unsubscribe him, so do it
        if(!$dbh->query('DELETE FROM newsletter_addresses
                         WHERE confirm_code = ?',array(md5($vcode)))){
                print "Error while accessing database,
                       contact system administrator."; die;
        }
}
print "TRUE";
?>
```

**Figure 5** The code of confirm.php, on which `confirm` in Fig. 1 is based.

## A.2 Correcting the Error in Panda Source Code

In Section 2 we described a natural correction of the error in the running example. Under this correction, `confirm − corrected` satisfies the pre- and postconditions $\text{pre}_{\texttt{confirm}}$ and $\text{post}_{\texttt{confirm}}$ from Fig. 2.

```
confirm-corrected(cd,sub_gh):
    A = SELECT subscribe FROM NS WHERE code = cd;
    if (A = empty) exit; //"No such code"
    if (sub_gh = false) UPDATE NS SET subscribed = true, code = nil WHERE code = cd
    else DELETE FROM NS WHERE code = cd;
```

$sub_{gh}$ is no longer a ghost variable. Now it is a second argument to `confirm − corrected`. The function `subscribe`, which had generated URLs calling `confirm` with one argument, namely the confirm code, now generates URLs with an additional argument *true*. Similarly, `unsubscribe` generates URLs with the additional argument *false*.

# B The Conference Management Example

In this example we verify parts of a system for conference management which assigns reviewers to papers and records the reviews and acceptance/rejection decisions. We focus on the earlier parts of the reviewing process: First, potential reviewers (e.g. PC members) bid on papers to review. Based on the bids, reviewers are assigned to the papers (e.g. by the PC chair). An additional functionality of the system that we focus on is displaying the list of papers by a specific author.

## B.1 The database

The database contains the following tables and columns:

**Papers** with columns *paperId*, *status*, and *session*. The column *status* is over the bounded domain consisting of *undecided*, *accepted*, or *rejected*. The column *session* ranges over the bounded domain consisting of *null*,*blank*,*invited*,1,. . .,*k*;

**PaperAuthor** with columns *userId* and *paperId*;

**ReviewerBids** with columns *userId* and *paperId*;

**ReviewerAssignments** with columns *userId* and *paperId*;

**Conflicts** with columns *userId* and *paperId*.

The columns *userId* and *paperId* range over the unbounded domain. The key of Papers is *paperId*. The other tables have a many to many relationship between *userId* and *paperId* attesting respectively to the fact that the user is the author of the paper, the user has bid to review the paper, the user has been assigned to review the paper, or the user is in conflict with the paper (and therefore cannot review it).

Before the bidding process begins, all papers are assigned the status *undecided* and the session *invited* (for an invited paper) or *null* (for a contributed submission). The session value *blank* comes up in the display function, and is at the root of a bug in the program.

## B.2 The functions `bid`, `assign`, and `display`

The functions `bid`, `assign`, and `display` are referred to as `conf-bid`, `conf-assign`, and `conf-display` in Table 1. The code of the functions `bid`, `assign`, and `display` can be found in Fig. 6. The function `bid` registers that the user *usr* is willing to review the paper *ppr* with the sanity check that there is no conflict between the user and the paper. The table $A$ is either empty whenever no conflict is found, or contains the single row *ppr* when there is a conflict. The function `assign` registers that the user *usr* is assigned to review the paper *ppr*. The function `display` receives as input the user id *usr* and returns the list of papers by *usr* that should be displayed. If the review phase of the conference is not yet completed (i.e. the Boolean argument *stillReviewing* has value true), `display` removes the

```
bid(usr, ppr):
    A = SELECT paperId FROM Papers WHERE paperId = ppr AND NOT (paperId IN
        (SELECT paperId FROM Conflicts WHERE userId = usr));
    if (A = empty) exit;
    INSERT (usr, ppr) INTO ReviewerBids


assign(usr,ppr):
    INSERT (usr, ppr) INTO ReviewerAssignments


display-incorrect(usr, stillReviewing):
    Output = SELECT * FROM Papers WHERE paperId IN
        (SELECT paperId FROM PaperAuthor WHERE userId = usr);
    if (stillReviewing = false) exit;
    UPDATE Output SET session = blank WHERE session IN 1,...,k


display-correct(usr, stillReviewing):
    Output = SELECT * FROM Papers WHERE paperId IN
        (SELECT paperId FROM PaperAuthor WHERE userId = usr);
    if (stillReviewing = false) exit;
    UPDATE Output SET session = blank WHERE session IN 1,...,k,null
```

**Figure 6** The code of `bid`, `assign`, `display-incorrect`, and `display-correct`.

session values of contributed papers from the output. This is done to prevent leaking the information that a contributed paper has been accepted (since only accepted papers have sessions) before the status of the paper has been announced. `display` leaves the status value *invited* visible.

We present two versions of `display`: one correct and one incorrect. `display-incorrect` leaves the session value *null* unchanged. Since *null* and *blank* are different values, the information leak which the program tries to avoid is still present. The correct version `display-correct` differs from `display-incorrect` by also replacing the status *null* by *blank*. This is done by expanding the WHERE condition of the UPDATE.

## B.3 The specification

The database preserves two invariants:

$$\begin{aligned} Inv_1 &= \forall_{\mathbf{d}} x, y. ReviewerBids(x,y) \rightarrow \neg Conflicts(x,y) \\ Inv_2 &= \forall_{\mathbf{d}} x, y. ReviewerAssignments(x,y) \rightarrow \neg Conflicts(x,y) \end{aligned}$$

These invariants state that no user may bid or be assigned to review a paper with which they are in conflict.

The specification of `display` is as follows:

$$\begin{aligned} \mathrm{pre}_{\texttt{display}} &= Inv_1 \wedge Inv_2 \\ \mathrm{post}_{\texttt{display}} &= Inv_1 \wedge Inv_2 \wedge \textit{no-leak} \\ \textit{no-leak} &= \textit{stillReviewing} \rightarrow (\forall_{\mathbf{d}} x, y.\ Output(x,y) \rightarrow (y = blank \vee y = invited)) \end{aligned}$$

This specification holds for `display-correct` and does not hold for `display-incorrect`.

For `bid` and `assign` we provide two specifications, one correct and one incorrect. The correct specification is as follows:

$$\begin{aligned} &\mathrm{pre}^c_{\texttt{bid}} = \mathrm{post}^c_{\texttt{bid}} = \mathrm{post}^c_{\texttt{assign}} = Inv_1 \wedge Inv_2 \\ &\mathrm{pre}^c_{\texttt{assign}} = Inv_1 \wedge Inv_2 \wedge ReviewerBids(usr, ppr) \end{aligned}$$

In order to ensure that $Inv_2$ is preserved by `assign`, we only allow a reviewer assignment to occur if there was a corresponding reviewer bid. Reviewer bids are required to avoid the conflicts by $Inv_1$,

and thus $Inv_2$ is preserved. The incorrect specification for `bid` and `assign` is as follows:

$$\begin{array}{lcl} \mathrm{pre}^{ic}_{\texttt{bid}} & = & Inv_2 \\ \mathrm{pre}^{ic}_{\texttt{assign}} & = & Inv_2 \wedge ReviewerBids(usr, ppr) \\ \mathrm{post}_{\texttt{bid}} & = & \mathrm{post}_{\texttt{assign}} = Inv_1 \wedge Inv_2 \end{array}$$

It is obtained from the correct specification by omitting $Inv_1$ from the preconditions.

# C The Firewall Example

In this example we verify a simple firewall with respect to a simple invariant. The firewall is provided with a database consisting of two tables: *Device* and *CanSend*. The table *Device* consists of a single column *deviceId*. The table *CanSend* consists of two columns *senderId* and *receiverId*. The *CanSend* table determines whether device A is allowed to send to device B.

The global invariant of the firewall requires that there exists a device to which every other device can send:

$$Inv_{\texttt{firewall}} = \exists x\,(Device(x) \wedge \forall y(Device(y) \rightarrow CanSend(y, x)))$$

We want to verify that this invariant holds when the network topology is changed. We consider the function `delete-device`:

```
delete-device(deviceToDelete):
    DELETE FROM CanSend WHERE senderId=deviceToDelete OR receiverId=deviceToDelete
    DELETE FROM Device WHERE deviceId=deviceToDelete
```

Table 1 refers to `delete-device` in the rows labeled *firewall*. In this table, we experiment with two specifications. The incorrect specification is:

$$\{Inv_{\texttt{firewall}}\}\ \texttt{delete} - \texttt{device}\ \{Inv_{\texttt{firewall}}\}$$

This specification is incorrect since it is possible that the only device which can receive messages from all other devices is exactly the device *deviceToDelete* removed by `delete-device`. Our correct specification is:

$$\{\mathrm{pre}_{\texttt{delete}-\texttt{device}} \wedge Inv_{\texttt{firewall}}\}\ \texttt{delete} - \texttt{device}\ \{Inv_{\texttt{firewall}}\}$$

where

$$\mathrm{pre}_{\texttt{delete}-\texttt{device}} = \exists x\,((x \neq deviceToDelete) \wedge Device(x) \wedge \forall y(Device(y) \rightarrow CanSend(y, x)))$$

This correct specification ensures that there is a device as required which is not *deviceToDelete*.

## C.1 Inflated examples

As a basic test of the scalability of our approach, in Section 5.3 we created large examples by inflating small examples. Here we illustrate this on the firewall example. The following is the result of inflating `delete-device` with multiplier 3. The function `delete-device3` deletes three devices from the network:

```
delete-device3(deviceToDelete1,deviceToDelete2,deviceToDelete3):
    DELETE FROM CanSend WHERE senderId=deviceToDelete1 OR receiverId=deviceToDelete1
    DELETE FROM CanSend WHERE senderId=deviceToDelete2 OR receiverId=deviceToDelete2
    DELETE FROM CanSend WHERE senderId=deviceToDelete3 OR receiverId=deviceToDelete3
    DELETE FROM Device WHERE deviceId=deviceToDelete1
    DELETE FROM Device WHERE deviceId=deviceToDelete2
    DELETE FROM Device WHERE deviceId=deviceToDelete3
```

The specifications must be altered correspondingly, so the precondition of the correct specification is changed to:

$$\begin{aligned}\text{pre}_{\texttt{delete-device3}} \quad = \quad & \exists x((x \neq \mathit{deviceToDelete1}) \wedge (x \neq \mathit{deviceToDelete2}) \wedge (x \neq \mathit{deviceToDelete3}) \\ & \wedge \mathit{Device}(x) \wedge \forall y(\mathit{Device}(y) \rightarrow \mathit{CanSend}(y, x)))\end{aligned}$$

The other pre- and postconditions remain unchanged, since they just consist of the invariant.

In Table 2, the rows for *firewall* correspond to the inflated versions of `delete-device` with the altered specifications.

# D FO$^2$ Reasoning

This appendix gives more detail on the implementation of our FO$^2$ finite satifiability solver. Section D.1 gives explicityly the Tseitin transformation of the $\forall\forall$-conjuct from Section 4.2.2. Section D.3 discusses our use of a refined bound on the size of the maximal model due to [21] to improve the efficiency of the solver.

## D.1 The $\forall\forall$-conjunct

Here we give $B^{\forall\forall}$ explicitly.

Let $\forall x \forall y\, \alpha'$ be the unique $\forall\forall$-conjunct of $\psi$. The CNF formula $B^{\forall\forall}$ will have clauses which explicitly detail the semantics of $\forall x \forall y\, \alpha'(x, y)$ on the structure $\mathcal{A}_\mathrm{S}$. We associate each sub-formula $\gamma(x, y)$ in $\alpha'(x, y)$ and $\ell_1, \ell_2 \in \{1, \ldots, \mathit{bnd}(\psi)\}$ with a new Boolean variable $u^{\gamma}_{\ell_1,\ell_2}$. $B^{\forall\forall}$ will contain clauses guaranteeing that under any assignment S satisfying $B_\psi$,

($\triangle$) $S(u^{\gamma}_{\ell_1,\ell_2}) = \mathit{True}$ iff $\mathcal{A}_\mathrm{S} \models \gamma(\ell_1, \ell_2)$.

Additionally, we add to $B^{\forall\forall}$ the clause $u^{\alpha'}_{\ell_1,\ell_2}$ for all $1 \leq \ell_1, \ell_2 \leq \mathit{bnd}(\psi)$ to assert that $\alpha'$ is true for all values of $x$ and $y$.

It remains to describe the clauses which define the values of $u^{\gamma}_{\ell_1,\ell_2}$ according to ($\triangle$). We distinguish two cases, depending on whether $\gamma$ is an atom of $\psi$ or is obtained by applying a Boolean connective $\vee, \wedge, \neg$ on sub-formulas. Consider first the case that $\gamma(x, y) = \neg(\delta(x, y))$. For every $1 \leq \ell_1, \ell_2 \leq \mathit{bnd}(\psi)$, we add the clauses

$$\neg u^{\gamma}_{\ell_1,\ell_2} \vee \neg u^{\delta}_{\ell_1,\ell_2} \quad \text{and} \quad u^{\gamma}_{\ell_1,\ell_2} \vee u^{\delta}_{\ell_1,\ell_2}$$

whose conjunction is equivalent to $u^{\gamma}_{\ell_1,\ell_2} \leftrightarrow \neg u^{\delta}_{\ell_1,\ell_2}$. The other Boolean connectives are axiomatized similarly.

For the case of atoms, $u^{\gamma}_{\ell_1,\ell_2}$ gets it value from one of the $v$ variables by choosing the indices correctly:

- If $\gamma(x, y) = R(x, y)$, then for every $1 \leq \ell_1, \ell_2 \leq \mathit{bnd}(\psi)$, $u^{R(x,y)}_{\ell_1,\ell_2}$ is assigned the same value as $v_{R,\ell_1,\ell_2}$ by adding the clauses

  $$\neg u^{R(x,y)}_{\ell_1,\ell_2} \vee v_{R,\ell_1,\ell_2} \text{ and } u^{R(x,y)}_{\ell_1,\ell_2} \vee \neg v_{R,\ell_1,\ell_2}$$

  to $B_\psi$. The conjunction of these two clauses is equivalent to $u^{R(x,y)}_{\ell_1,\ell_2} \leftrightarrow v_{R,\ell_1,\ell_2}$.
- If $\gamma(x, y) = R(y, x)$, then $u^{R(x,y)}_{\ell_1,\ell_2}$ is assigned the same value as $v_{R,\ell_2,\ell_1}$.
- If $\gamma(x, y) = R(x, x)$, then $u^{R(x,x)}_{\ell_1,\ell_2}$ is assigned the same value as $v_{R,\ell_1,\ell_1}$.
- If $\gamma(x, y) = R(y, y)$, then $u^{R(y,y)}_{\ell_1,\ell_2}$ is assigned the same value as $v_{R,\ell_2,\ell_2}$.
- If $\gamma(x, y) = U(y)$, then $u^{U(y)}_{\ell_1,\ell_2}$ is assigned the same value as $v_{U,\ell_2}$.
- If $\gamma(x, y) = U(x)$, then $u^{U(x)}_{\ell_1,\ell_2}$ is assigned the same value as $v_{U,\ell_1}$.

## D.2 Axiomatizing models of at size at most $bnd(\psi)$

Here we continue the discussion postponed to the appendix in Section 4.2.2. Recall that by Proposition 2, $B_\psi$ is satisfiable iff $\psi$ is satisfiable by a structure of cardinality *exactly* $bnd(\psi)$. However, [21] guarantees only that $bnd(\psi)$ is an *upper bound* on the cardinality of a satisfying model. In this appendix we explain how to construct $C_\psi$ so that it is satisfiable iff $\psi$ is satisfiable by a structure of cardinality *at most* $bnd(\psi)$ as follows.

We compute from $\psi$ a new $\mathrm{FO}^2$-sentence $\mathrm{Uni}(\psi)$ in SSNF and set $C_\psi = B_{\mathrm{Uni}(\psi)}$. Let *Uni* be a fresh unary relation name. Let $\mathrm{Uni}(\psi)$ be:

$$\forall x \forall y \left( \mathrm{Uni}_\alpha(x,y) \wedge \bigwedge_{i=1}^{m+1} F_i(x,y) \rightarrow \mathrm{Uni}_{\beta_i}(x,y) \right) \wedge \bigwedge_{i=1}^{m+1} \forall x \exists y\, F_i(x,y)$$

with $\mathrm{Uni}_\alpha(x,y) = \mathit{Uni}(x) \wedge \mathit{Uni}(y) \rightarrow \alpha(x,y)$, $\mathrm{Uni}_{\beta_i}(x,y) = \mathit{Uni}(x) \wedge \mathit{Uni}(y) \rightarrow \beta_i(x,y)$, $i = 1, \ldots, m$, and $\mathrm{Uni}_{\beta_{m+1}}(x,y) = \mathit{Uni}(y)$. $F_{m+1}$ is used to guarantee that *Uni* is non-empty. Let $\mathrm{voc}(\mathrm{Uni}(\psi))$ be the vocabulary of $\mathrm{Uni}(\psi)$.

▶ Proposition 3. Let $\mathcal{A}$ be a $\mathrm{voc}(\mathrm{Uni}(\psi))$-structure. Let $\mathcal{A}_{\mathrm{Uni}}$ be the substructure of $\mathcal{A}$ whose universe is $\mathit{Uni}^{\mathcal{A}}$. We have $\mathcal{A} \models \mathrm{Uni}(\psi)$ iff $\mathcal{A}_{\mathrm{Uni}} \models \psi$.

## D.3 An efficient finite satisfiability algorithm

The algorithm from Section 4.1 on which our algorithm in Section 4.2 is based was written from a theoretical point of view aiming to simplify the proof of the NEXPTIME-completeness of the finite satisfiability problem. In this section, we introduce several optimizations which, while not affecting the complexity of the problem, improve the performance of our satisfiability solver.

### D.3.1 A refined upper bound

In the course of the proof of the bounded model property, [21] give a more refined version of the upper bound on the size of a minimal satisfying model. This more refine version leads to smaller upper bounds in many cases. To state the refined upper bound we need some definitions.

▶ **Definition 2. (1-types and kings)**

1. A 1-type $t(x)$ is a maximally consistent set of atomic formulas and their negations which do not have $y$ as a free variable.
2. For a structure $\mathcal{A}$ and an element $a$ of the universe of $\mathcal{A}$, the 1-type of $a$ in $\mathcal{A}$ is the unique 1-type $t(x)$ such that $\mathcal{A} \models t(a)$. We say that $a$ *realizes* $t$ *in* $\mathcal{A}$.
3. Given a structure $\mathcal{A}$ and an element $a$ of $\mathcal{A}$, $a$ is a *king* in $\mathcal{A}$ if there is no other element in $\mathcal{A}$ with the same 1-type as $a$.

For example, for a vocabulary consisting of one binary relation name $R$ and one constant name $c$, the following is a 1-type:

$$\{R(c,c), R(x,c), \neg R(c,x), R(x,x),\\ \neg(x=c), \neg(c=x), (x=x), (c=c)\}$$

▶ **Lemma 3** ([21], Theorem 4.3). *Let $\psi$ be a sentence in Skolemized Scott Normal Form and let $m$ be the number of conjuncts of the form $\forall x \exists y\, \beta_i$ as in Eq. (2). Let $\mathcal{A}$ be a structure satisfying $\psi$. Let $K$ be the set of all kings in $\mathcal{A}$ and let $P$ be the set of 1-types realized in $\mathcal{A}$. There is a structure of cardinality at most*

$$(m+1)|K| + 3m(|P| - |K|) \tag{3}$$

*which satisfies $\mathcal{A}$.*

The bound in Lemma 3 requires already having a model of $\psi$. However, we can use it to get a bound based on syntactic considerations only. For $m \geq 1$, Eq. (3) can be bounded from above by the sum over all 1-types $t(x)$, such that $t(x)$ contributes $(m+1)$ if $t(x)$ contains $x = c$ for some $c \in \mathrm{const}(\psi)$,

and $3m$ otherwise. Note that the number of 1-types containing $x = c$ for some $c \in \text{const}(\psi)$ in any one structure is at most $|\text{const}(\psi)|$. Hence, Eq. (3) is at most:

$$|\text{const}(\psi)|\,(m+1) + \sum_{t(x) \not\models \bigvee_{c \in \text{const}(\psi)} x=c} 3m \tag{4}$$

Notice that this bound does not depend on whether the 1-types in the sum are realized in any structure. It is correct since (1) any 1-type which implies that $x = c$ for any constant is necessarily a king and thus contributes $m + 1$, while (2) any other 1-type, which may or may not be a king, contributes at most $\max(m + 1, 3m) = 3m$.

We can now augment the algorithm in Section 4.2 to use the refined bound from Eq. (4). Our algorithm uses Proposition 1 to transform a FO$^2$-sentence $\phi$ into Scott Normal Form, which adds a new relation name for every quantifier and every connective in $\phi$. This comes at a heavy cost to performance, since the number of 1-types summed over in Eq. (4) is exponential in the number of relation and constant names in Scott Normal Form of $\phi$. In this section we provide a more economic procedure for this purpose, which introduces new relation names only as a last resort.

Given a FO$^2$-sentence $\phi$, we construct a sentence $\psi$ in Scott Normal Form such that $\phi$ and $\psi$ are satisfiable by models of the same cardinalities. We construct a sequence of pairs $(\phi_k, \psi_k)$, $1 \leq k \leq r$ as follows. The sequence is built according to the process described below. The length $r$ of the sequence is determined by applying the process until no further steps can be applied. The sequence satisfies:

- $\phi_0$ is the Negation Normal Form[6] of $\phi$ and $\psi_0 = \forall x \forall y\, True$,
- $\phi_r = True$ and $\psi_r = \psi$,
- $\psi_k$ is in Scott Normal Form for all $1 \leq k \leq r$, and
- the sets of cardinalities of the models of $\phi_k \wedge \psi_k$ are equal for all $1 \leq k \leq r$.

Given $(\phi_k, \psi_k)$ we compute $(\phi_{k+1}, \psi_{k+1})$ iteratively as follows:

**1.** If $\phi_k$ is in one of the forms:

$$\forall x Q y\, \epsilon \qquad \forall x\, \epsilon$$
$$\forall y Q x\, \epsilon \qquad \forall y\, \epsilon$$

where $\epsilon$ is quantifier-free and $Q$ is a quantifier, i.e. $Q \in \{\exists, \forall\}$, then $\psi_{k+1}$ is obtained from $\psi_k$ as follows. If $Q = \forall$, $\psi_{k+1}$ is obtained by adding $\epsilon$ as a new conjunct inside the quantifiers $\forall x \forall y$. If $Q = \exists$, $\psi_{k+1}$ is obtained by adding a new conjunct $\forall x \exists y\, \epsilon$ or $\forall y \exists x\, \epsilon$ to $\psi_k$. We set $\phi_{k+1}$ to True. We end the iteration by setting $r$ to $k + 1$,

**2.** If $\phi_k$ is a conjunction in which one of the conjuncts $\gamma$ is of one of the forms in the previous item, then $\psi_{k+1}$ is obtained from $\psi_k$ as in the previous item and $\phi_{k+1}$ is obtained from $\phi_k$ by removing $\gamma$ from the conjunction.

**3.** If there is an existential quantifier not in the scope of any other quantifier in $\phi_k$, then $\phi_{k+1}$ is obtained by removing the existential quantifier and replacing all the occurrences of the quantified variable bound to this quantifier with one fresh constant name. We set $\psi_{k+1} = \psi_k$. Note for the correctness of this step that we are using here that the formulas are in Negation Normal Form, i.e. this existential quantifier may be in the scope of the $\wedge$ and $\vee$ operators only.

**4.** If $\phi_k$ is of one of the following forms, or if $\phi_k$ is a conjunction in which one of the conjuncts $\gamma$ is of one of the following forms:

$$\forall z_1\, ((Q z_2\, \delta_1) \vee \delta_2)$$
$$\forall z_1\, ((Q z_2\, \delta_1) \wedge \delta_2)$$
$$\forall z_1\, (\delta_1 \vee (Q z_2\, \delta_2))$$
$$\forall z_1\, (\delta_1 \wedge (Q z_2\, \delta_2))$$

[6] An FO formula is in Negation Normal Form (NNF) if the scope of every negation symbol $\neg$ is an atom. It is well-known that for every FO formula $\gamma$, an equivalent formula $\gamma'$ in NNF can be computed in linear time.

where $Q \in \{\exists, \forall\}$, $z_1, z_2 \in \{x, y\}$ and $z_1 \neq z_2$, $\phi_{k+1}$ is obtained by taking the quantifier $Qz_2$ out of the scope of the Boolean connective. E.g., we substitute $\forall z_1((Qz_2\delta_1) \vee \delta_2)$ with $\forall z_1 Qz_2(\delta_1 \vee \delta_2)$. We set $\psi_{k+1} = \psi_k$. Note for the correctness of this step that the $\delta_j$ not in the scope of $Q$ does not have $z_2$ as a free variable, since we are dealing with sentences.

5. If $\phi_k$ is quantifier-free, $\psi_{k+1}$ is obtained from $\psi_k$ by adding $\phi_k$ as a new conjunct inside the quantifiers $\forall x \forall y$ in $\psi_k$. We set $\phi_{k+1}$ to True and end the process by setting $r = k + 1$.
6. If none of the previous items applied to $\phi_k$ in this iteration, we eliminate one quantifier from $\phi_k$ in spirit of the discussion before Proposition 3.1 in [21]. Let $\gamma$ be a sub-formula of $\phi_k$ of the form $Qz\,\delta$, where $Q \in \{\exists, \forall\}$, $z \in \{x, y\}$ and $\delta$ is quantifier-free. Let $\bar{z} \in \{x, y\}$ such that $z \neq \bar{z}$. The sentence $\phi_{k+1}$ is obtained by substituting $\gamma$ by $E(\bar{z})$ in $\phi_k$, where $E$ is a fresh unary relation name. Let $\theta = \forall \bar{z}\,((Qz\delta) \leftrightarrow E(\bar{z}))$. The sentence $\theta$ says that $E$ is interpreted as the set of universe elements $u$ for which $Qz\delta$ holds. $\theta$ is equivalent to the conjunction of $\theta_{\forall\forall}$ and $\theta_{\forall\exists}$ such that $\theta_{\forall\forall}$ is of the form $\forall x \forall y\, \theta'_{\forall\forall}$ and $\theta_{\forall\exists}$ is of the form $\forall x \exists y\, \theta'_{\forall\exists}$, and $\theta'_{\forall\forall}$ and $\theta'_{\forall\exists}$ are quantifier-free. Let $\psi_{k+1}$ be obtained by adding the conjunct $\theta_{\forall\exists}$ to $\phi_k$ and adding $\theta'_{\forall\forall}$ as a new conjunct inside the $\forall\forall$-conjunct of $\phi_k$.

Only item 6 increases the number of names, so it is only used when no other item applies. The procedure terminates because every item, except for item 4, removes a quantifier when going from $\phi_k$ to $\phi_{k+1}$, and whenever item 4 is applied, in the next iteration either item 1 or item 2 will be applied.

### D.3.2 Ruling out unfeasible 1-types

Up until now, we have bounded the number of 1-types which are realized in some structure $\mathcal{A}$ satisfying $\phi$ with the number of all 1-types. However, it is possible to determine that some 1-types are not feasible in any structure satisfying $\phi$ and subtract them from the upper bound.

Recall the table *NS* from the running example in Section 2. $NS(x, y)$ expresses that user $y$ is subscribed to newsletter $x$. It is natural that the requirement that users and newsletters are disjoint

$$\forall x\ ((\exists y\, NS(x, y)) \rightarrow (\forall y\, \neg NS(y, x)))$$

is part of the database invariant. Hence, any 1-type containing $NS(x, x)$ is unfeasible.

Let $\psi$ be a sentence in Scott Normal Form such that $\psi = \psi_{\forall\forall} \wedge \psi_{\forall\exists}$, where $\psi_{\forall\forall} = (\forall x \forall y\, \alpha)$ and $\psi_{\forall\exists}$ is a conjunction of terms of the form $\forall x \exists y\, \beta_i$. Any 1-type which is not feasible for $\psi_{\forall\forall}$ is certainly not feasible for $\psi$. Since $\psi_{\forall\forall}$ is a universal FO sentence, it adheres to a classical property of universal FO: the class of models of $\psi_{\forall\forall}$ is closed under taking substructures. This implies that any feasible 1-type of $\psi_{\forall\forall}$ occurs in a structure of cardinality 1. We count the number of unfeasible 1-types which do not contain $x = c$ for any constant name by iterating over all such 1-types $t(x)$ and checking satisfiability of $\forall x\, t(x) \wedge \psi_{\forall\forall}$ by a structure of cardinality 1.

### D.3.3 Improved satisfiability testing

The size of the input and the run time of the SAT solver depend on the size of the upper bound $bnd(\psi)$ on the cardinality of the minimal satisfying model, if such a model exists. To guarantee that $\psi$ is unsatisfiable, one must verify that no satisfying model up to that cardinality exists. In contrast, to guarantee that $\psi$ is satisfiable, it is enough to find a satisfying model, which may be much smaller than the bound. Hence, instead of performing the satisfiability test for models of cardinality up to $bnd(\psi)$, we iteratively search for models whose cardinalities $1 = a_0, \ldots, a_e \leq bnd(\psi)$ increase exponentially up to $bnd(\psi)$. For each $a_i$, we construct a CNF formula $C_{a_i}$ similar to $C_\psi$, only replacing the maximal cardinality $bnd(\psi)$ with $a_i$, and apply the SAT solver to $C_{a_i}$. We have that $a_i = 2a_{i-1}$ for every $2 \leq i < e$ and $a_e = \min(bnd(\psi), 2a_{e-1})$. The search continues until either we reach the index $e$ for which $a_e = bnd(\psi)$ or $C_{a_e}$ is satisfiable, and the algorithm returns the truth-value of $C_{a_e}$. This procedure is very often much faster for satisfiable inputs $\psi$. For unsatisfiable inputs, the run time increase is negligible, since the formulas $C_{a_0}, \ldots, C_{a_{e-1}}$ are exponentially smaller than $C_{a_e} = C_\psi$.